

Adaptable Assignment

William G. Allen, Jr.

Abstract

This paper reports on a practical, simple method for adjusting a vehicle trip table so that the resulting assignments more closely match available traffic counts. “Practical” means that this is not purely a research effort — the procedure described here has been used in several studies as production-ready tool. “Simple” means that no special programming is necessary — the method can be implemented using existing travel modeling software packages. This procedure does not create a trip table directly from only traffic counts, but rather adjusts an existing table.

Over the years, researchers have developed a number of procedures which claim to synthesize travel patterns (i.e., a trip table) more or less directly from a set of traffic counts. This has great appeal to transportation planners, who are often faced with the task of quickly developing a traffic forecasting procedure based on little except a zone system, a network, and some counts (and sometimes, not even that much). However, most of these methods have not been successfully transferred to real world practitioners. Those few techniques that are in use are often viewed as “black boxes” that use obscure, complex algorithms and/or specialized software.

The method discussed here avoids elaborate algorithms in favor of a more easily understood “brute force” approach. The process requires the construction of minimum paths and the summation of the total count and total assigned volume (on links with counts) for each O-D pair. This information can readily be used to adjust the trips in each O-D pair. By iterating the process many times, dramatic improvements in the root-mean square assignment error can be achieved without unduly distorting the original trip table. This has been implemented as a simple multi-step procedure using the MINUTP software package; it may be possible to apply the procedure using other packages as well. The process results in a “delta” trip table, representing the difference between the original and adjusted trip tables. This delta table can be used by itself in forecasting or as the basis from which a model can be adjusted.

This procedure has been used in several cities in New Jersey and Pennsylvania, with satisfactory results.

This paper reports on a practical, simple method for adjusting an existing vehicle trip table so that the resulting assignments more closely match available traffic counts. “Practical” means that this is not purely a research effort — the procedure described here has been used in several studies as a production-ready tool. “Simple” means that no special programming is necessary — the method can be implemented using an existing travel modelling software package. This procedure does not create a trip table directly from only traffic counts, but rather adjusts an existing table.

Over the years, researchers have developed a number of procedures which claim to synthesize travel patterns (i.e., a trip table) more or less directly from a set of traffic counts. This has great appeal to transportation planners, who are often faced with the task of quickly developing an assignment procedure based on little except a zone system, a network, and some counts (and sometimes, not even that much). However, most of these methods have not been successfully transferred to real world practitioners. Those few techniques that are in use are often viewed as “black boxes” that may use obscure, complex algorithms and/or specialized software.

For example, The Highway Emulator (THE) software package incorporates a “trip table from counts” methodology. TMODEL2 also includes this feature, using an algorithm based on Wil-lumsen’s Method. Cambridge Systematics has developed a program it calls “TTE” for this purpose. These methods all use some variation of an iterative proportional fitting technique. The TRIPS software package uses a sophisticated mathematical programming technique in its MVESTM procedure to estimate trip tables. The EMME/2 and TransCAD packages reportedly also include such procedures.

The method described here, called *adaptable assignment*, takes a different approach. Instead of elaborate algorithms, it uses a very simple “brute force” set of calculations. Instead of a black box compiled program, it uses a straightforward set of modelling steps, which any reasonably capable software package should be able to perform.

One of the dangers of *any* kind of trip table estimation process is that it can lead to distorted and unrealistic trip patterns, depending on the count coverage and the starting trip table. With adaptable assignment, the process is simple and open and the analyst can easily check the results at any point in the process. This makes it much easier for the analyst to ensure that the resulting trip table is not unreasonable. Also, since the process can be applied using standard software, the extra cost of specialized programs is avoided. With continuing advances in microprocessor speed, brute force methods become increasingly feasible.

Methodology

Adaptable assignment is an iterative procedure by which an assignment is made, the resulting assigned volumes are systematically compared to the counts, trip table adjustment factors are developed and applied, and the process repeated. As developed thus far, there is no means of automatically stopping the process based on some criterion of convergence, but it should be possible to add that feature at a later date. Thus, the analyst must specify a fixed number of iterations and include a step to calculate and display whatever convergence statistic is desired.

As described here, the procedure has been implemented using the MINUTP software package. The process uses one special feature of MINUTP, but it may be possible to implement the process within other packages, such as TRANPLAN and EMME/2. The process is implemented as a DOS batch file, which runs a series of four MINUTP steps through several iterations. Thus far, it appears that 10-20 iterations provides an acceptable trade-off between accuracy and processing time. Continuing advances in microprocessor speed make it increasingly feasible to run numerous iterations of this procedure, and the process can be entirely automated via batch files and can run unattended.

The four main steps of the procedure are as follows:

Step 1: Traffic Assignment

The analyst must begin with a highway network and a starting vehicle trip table. The trip table can represent daily or peak hour trips, but must be consistent with the counts that are posted in the network. The trip table can come from any source: calibrated gravity model, synthesized gravity model, previous year’s trip table, travel survey, etc. Obviously, the more confidence the analyst has in the starting trip table, the better the results. However, it is believed that even a poorly synthesized trip table based on estimated land use, approximate trip rates, and borrowed F factors, should be an acceptable starting point.

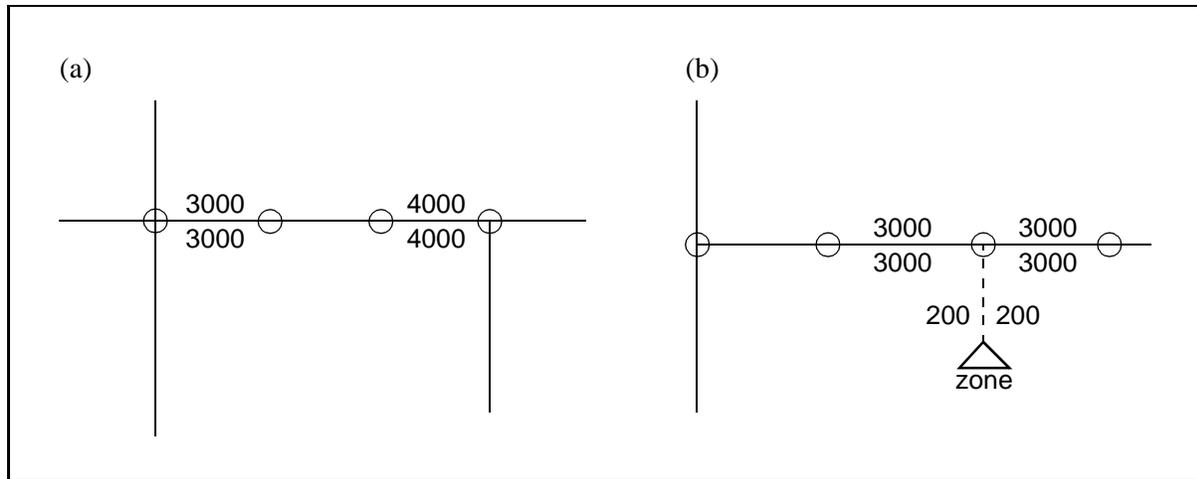


Figure 1: Inconsistent counts

The network must contain *some* counts. The networks to which this process has been applied thus far have contained counts on 7-8% of their links. It is believed that the process will work with as little as 4% count coverage, but obviously, the more counts, the better. Counts should be directional, for ease of processing. If daily non-directional counts are available, they should be divided in half and posted on each side of the road.

It is extremely important that the count data be internally consistent and logical. Situations such as those shown in Figures 1a and 1b must be avoided. This is because regardless of the accuracy of the count, it is impossible for assigned volumes to *ever* match the counts shown in these figures. In such cases, counts could be averaged, or one count could be removed in favor of the other. Inconsistent counts will distort the process, leading to unreliable results. This is true for *any* assignment validation process. (Actually, the situation in Figure 1b is technically possible, but extremely unlikely.)

Any assignment procedure can be used: equilibrium, incremental (CATS), all-or-nothing, stochastic, etc. The MINUTP ASSIGN program can accommodate a variety of procedures. As developed thus far, the adaptable assignment process assumes that the assignment procedure is fixed and that only the trip table varies. The output of Step 1 is a loaded network which contains both assigned volumes and counts.

Step 2: Network Calculations

MINUTP provides an all-purpose network calculator called NETMRG. In this step, NETMRG is used for three things:

2a) Calculate Assignment Accuracy

The square of the error $(\text{count} - \text{assigned})^2$ is calculated for each link with a count and summed for the network. If the number of links with a count is known, the total squared error can be used to manually calculate the percent root-mean-square error (%RMSE), which is widely used to measure assignment accuracy. (The most recent version of NETMRG calculates the RMSE automatically.) The %RMSE is calculated as follows:

$$\%RMSE = 100 \times \frac{\sqrt{\sum (count - assignment)^2}}{\sum count}$$

n

where n = number of links with a non-zero count (in some formulations, the n term inside the square root is replaced with n-1)

2b) Compute “Assignment Criteria”

For reasons that are explained below, two new network variables are calculated for each link: the count multiplied by a “scaling factor” (a fraction) and the assigned volume multiplied by the same fraction (computed only for links with a non-zero count). A typical value for this scaling factor is 0.1.

2c) Check Major Links

The count, assigned volume, and percent error for major links are calculated and listed. “Major” is usually defined as links above a certain count volume or links of a certain facility type, but any similar criterion can be used.

The result of Step 2 is another loaded network, with two new variables, plus a listing file that can be used to check the procedure’s progress.

Step 3: Skim Assignment Criteria

This is the heart of the procedure. Almost all software packages can “skim” a network’s time and distance (i.e., compute minimum paths and sum the time and distance along those paths for each origin-destination (O-D) pair). MINUTP’s PTHBLD program can skim almost any variable and that feature is used here to skim the two assignment criteria: the factored count and the factored assignment. The only limitation (in MINUTP) is that for any O-D pair, the total skimmed count and the total skimmed assigned volume must each be 32,767 or less. This is why the assignment criteria are computed in Step 2b as a fraction of the original count and assigned volume.

The result of this step is two matrices: one representing a fraction of the total counted volume on the minimum path for each O-D pair and another representing the same fraction of the assigned volume on counted links on the minimum path for each O-D pair.

Step 4: Trip Table Adjustment

The final step uses MINUTP’s MATRIX program to calculate an adjustment factor for each O-D pair with non-zero trips (this process does not insert values into zero trip cells). The ratio of the two matrices from Step 3 is: total factored count/total factored assignment. This ratio indicates the direction and magnitude of the “error” for each O-D pair.

It appeared desirable to provide for some damping of the calculation so as to prevent undesirable oscillation of the resulting estimates. Thus, a “sensitivity factor” is introduced, to dampen the adjustment, resulting in the following equation for the adjustment factor:

$$adjustment\ factor_{O-D} = \left(\frac{\sum factored\ count_{O-D}}{\sum factored\ assignment_{O-D}} \right)^{SF}$$

where SF = sensitivity factor

Thus far, reasonable results have been obtained with a sensitivity factor of 0.50. Higher values will result in faster convergence, but at the possible expense of undue trip table distortion. The “optimal” value of this parameter, if one exists, must await further research.

Once the adjustment factor is known, the new trips are calculated as:

$$\text{new trips}_{O-D} = \text{adjustment factor}_{O-D} * \text{old trips}_{O-D}$$

In this calculation, extra care must be taken to ensure the integrity of the trip table cell values. Bucket rounding or some other technique should be used to ensure that the resulting integer trip values reflect the intended calculation as accurately as possible. Finally, the delta trip matrix is calculated as:

$$\text{delta trips}_{O-D} = \text{new trips}_{O-D} - \text{old trips}_{O-D}$$

MINUTP is capable of handling and storing negative trip table entries, although special care must be used.

These four steps constitute one iteration of the procedure. It is necessary to repeat Steps 1 and 2a to see how well the adjusted trip table assigns. In practice, the entire procedure is applied 10-20 times, with the adjusted trip table from each iteration used as the input for the next iteration. Examination of the %RMSE value that is derived from Step 2a provides indicates how well the process is working.

One feature which can be added to the basic process is a FRATAR step after, say, every fifth iteration. The purpose of this is to force the trip ends of the adjusted trip table to be equal to some desired values, usually the original trip ends. In some cases, the original trip ends are deemed to be acceptable and the analyst would want the process to change only the individual cell values, while maintaining the original row and column totals. Frataring the adjusted trip table back to the original trip ends accomplishes this (approximately). By applying the FRATAR process after every fifth iteration, the analyst ensures that the adjusted trip table never strays too far from the original. It should be clear, however, that the FRATAR process disturbs the adjustments made by this procedure and that a less accurate assignment will result. In some instances, this may be an acceptable trade-off.

Another optional feature is that it is possible to adjust the trip table so that the assignment is extremely accurate for a select set of links. In some cases, the analyst may need for the assignments to match the counts on a small set of links to a very high degree of accuracy. This can be accomplished by adding a dummy variable to the network, named something like PRIORITY. The selected links are given a PRIORITY of 1 and all other links are given a PRIORITY of zero. Then, the equations for the assignment criteria become:

$$\text{factored count} = \text{original count} * 0.1 * \text{PRIORITY}$$

$$\text{factored assigned volume} = \text{original assigned volume} * 0.1 * \text{PRIORITY}$$

Essentially, this is equivalent to posting counts only on those links in the selected set. If there are fewer than, say, ten links in the selected set, applying adaptable assignment in this manner will practically guarantee that the assigned volumes will match the counts on those links. This process has also proven useful in trying to match counts that are very different in magnitude (say, daily traffic of 1,000 and 50,000 on nearby links).

It must be recognized that this process can modify the starting trip table in undesirable ways. Thus, the analyst must carefully examine the adjusted trip table. It is necessary to analyze the trip length distributions of the original and revised tables as well as district-to-district trip patterns, to be assured that the adjustments are appropriate.

Forecasting

The above discussion describes how to obtain an adjusted trip table, or alternatively, a delta trip table. The more interesting question is what to do with that table. Some analysts believe that it is appropriate to simply carry the delta trip table along into forecasting. For example, the analyst would use some procedure to obtain a future year trip table and would then add the delta trip table to it to obtain the final future trips.

Other analysts feel it is appropriate to develop a base year “ratio” table, in which the cells represent the ratio of the final adjusted trip table to the original trip table, factored by 1000 to represent integer values (most packages store matrix values as integers, so a ratio of 2.145 would need to be stored as 2145, with an understood divisor of 1000). The original future year table would be factored by this ratio table to obtain the final future trips.

While either of these approaches might be acceptable in some studies, it is unclear if they would be found acceptable by a “best practices” or peer review. If resources and circumstances permit, it seems more appropriate to use the delta trip table to modify the procedure used to obtain the original trip table. This could involve, for example, changes to the trip generation rates, identification of special generator zones, use of K factors in distribution, and mode share adjustments. In effect, the delta trip values would be integrated back into the model itself.

Sample Application

The adaptable assignment procedure has been used in Princeton, NJ, Reading, PA, the US 1 corridor in central New Jersey, and West Windsor Township, NJ. The West Windsor case is the best documented and is described here. West Windsor is a small town located in the heavily congested US 1 corridor north of Princeton, close to the geographic center of the state. The network consists of 275 zones and 1,600 links. The trip table represents 1993 peak hour vehicle trips and was synthesized from a combination of approximate trip rates, borrowed F factors, and cordon counts. About 7% of the links had posted counts. The following parameters were used:

- scaling factor = 0.1
- sensitivity factor = 0.50
- iterations = 20
- no Fratarling of trip ends

Table 1 shows the AM and PM results and Figure 2 shows the PM change in %RMSE by iteration. As these results show, the improvement in assignment accuracy is substantial. This run

Table 1: West Windsor results

	AM	PM
Total Count	89,152	118,076
Initial Total Assigned Volume	103,687	147,325
Final Total Assigned Volume	89,198	117,521
Total Volume Error	+0.05%	-0.47%
Initial Trips	20,324	23,995
Final Trips	19,264	21,674
Total Delta	-1,060	-2,321
Pct. Difference	-5.2%	-9.7%
Initial %RMSE	42.7%	48.2%
Final %RMSE	21.9%	16.2%

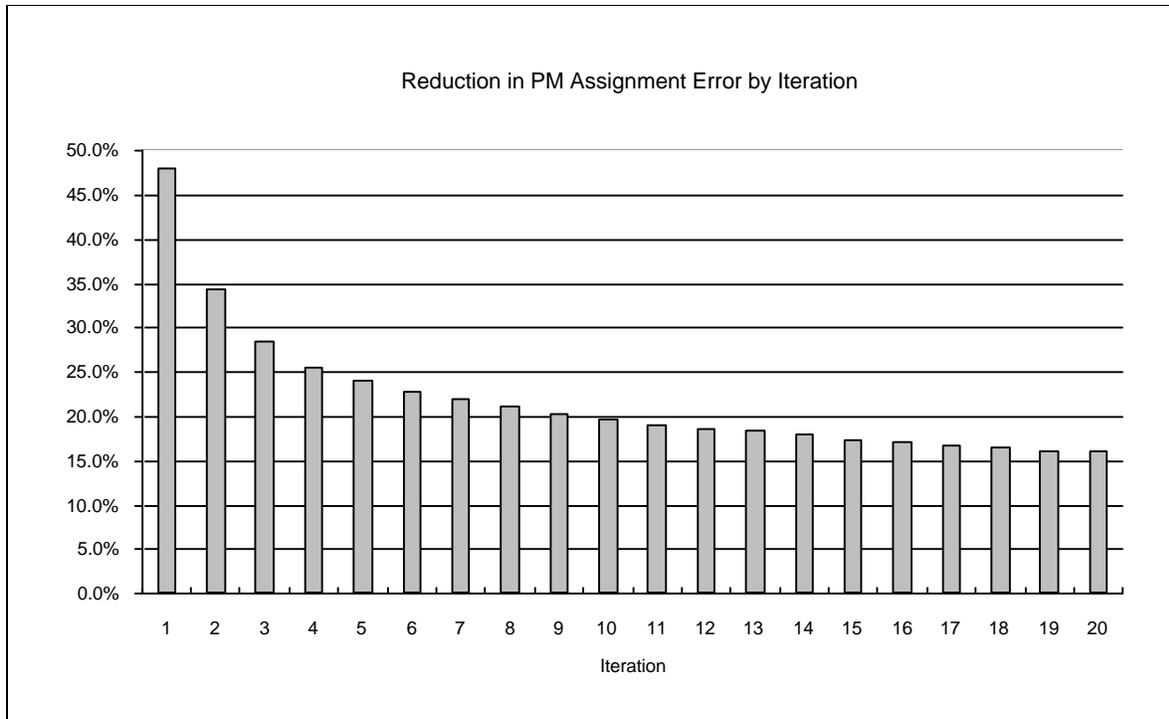


Figure 2: Reduction in p.m. assignment error by iteration

could probably have been stopped after 10 iterations, since the remaining improvement is fairly small. In fact, the overall improvement in accuracy is probably better than what can generally be expected, because the study area is a fairly small network. As Table 1 shows, the changes to the original trip table are fairly minor, in total.

Conclusions

A new procedure has been developed to adjust a trip table to better match available traffic counts. This “adaptable assignment” procedure has the advantages of being simpler and more accessible to users than other techniques. Although the procedure was developed using MINUTP, it might be possible to apply it using other software packages, but in any case, no expensive software or compiled programs are needed.

This procedure has been applied in a number of real-world studies and has been shown to produce substantial improvements in assignment accuracy without unduly distorting the original trip table. It is believed that by simplifying the process of trip table adjustment, adaptable assignment will enable more transportation planners to use and understand this tool.

Acknowledgments

The adaptable assignment procedure was developed by the author with the assistance and advice of Gary Davies of Garmen Associates. Most of the case studies mentioned here were studies conducted by Garmen. The opinions expressed are those of the author, who bears sole responsibility for the content of this paper.

Bibliography

Short-Term Travel Model Improvements, prepared by Cambridge Systematics Inc. for US DOT, October 1994.

TMODEL2 User's Manual, TMODEL Corporation, August 1994.

MINUTP Technical User's Manual, COMSIS Corporation, January 1996.